Network Working Group
Request for Comments: 1124

B. Leiner
RIACS

September 1989

Policy Issues in Interconnecting Networks

Status of this Memo

To support the activities of the Federal Research Internet Coordinating Committee (FRICC) in creating an interconnected set of networks to serve the research community, two workshops were held to address the technical support of policy issues that arise when interconnecting such networks. Held under the auspices of the Internet Activities Board at the request of the FRICC, and sponsored by NASA through RIACS, the workshops addressed the required and feasible technologies and architectures that could be used to satisfy the desired policies for interconnection.

The purpose of this RFC is to report the results of these workshops. Distribution of this memo is unlimited.

Preface

This report documents the results of two workshops held at the request of the Federal Research Internet Coordinating Committee and under the auspices of the Internet Activities Board. As such, this report represents the work of a large number of people (listed in Section 7). Without their efforts, these results would not have been possible. The author (really more of an editor) would like to acknowledge their efforts and contributions, and thank them for their cooperation in making the workshops a success.

1. Introduction

Computer networking has become pervasive and basic to the conduct of scientific and academic activities. To provide the needed networking support to these activities, each of the agencies funding research has proceeded to establish one or more agency funded computer networks.

Recognizing the importance of such networking support, the Office of Science and Technology Policy (OSTP) working with the appropriate personnel from the research-funding agencies on the Federal Coordinating Council on Science Engineering and Technology (FCCSET) Committee on High-Speed Networks developed a set of recommendations for the evolution and enhancements of scientific and academic networks. These recommendations are described in three phases. The first phase addresses the interconnection of the various agency networks into a ubiquitous networking capability serving several hundred universities and research institutions with a backbone network operating 1.5 Mb/s. The second phase involves upgrading the network backbone to 45 Mb/s and connecting additional universities and other research institutions. The third phase involves the development and installation of a high bandwidth (Gb/s) networking capability.

The motivation for the first two phases are to achieve good performance in a cost effective manner. The scientific and academic community is best served by an interconnected ubiquitous networking capability rather than a set of partitioned networks supporting only subsets of the community. Costs can be reduced and performance improved through sharing of resources and using cross-support (e.g., using one agency's network to serve an institution for another agency's purposes rather than having to connect each institution to every network.)

To accomplish these objectives, the Federal Research Internet Coordinating Committee (FRICC) was formed. Consisting of representatives from the key research agencies (NSF, DARPA, NASA, and DOE), this ad hoc group has been developing strategies for interconnection of networks and evolution of the Internet in accordance with the OSTP recommendations for Phases 1-3. In the process of developing such plans, it became apparent that a set of issues needed to be addressed concerning the various agency policies for their research networks in light of the desire to interconnect such networks.

This report documents the results of a series of two workshops (18-20 June 1988 at NASA Ames Research Center and 8-10 November 1988 at MIT) held to address these issues. Held under the auspices of the Internet Activities Board (IAB) at the request of the FRICC, and sponsored by NASA through RIACS, the workshops addressed the required and feasible technologies and architectures that could be used to satisfy the desired policies for interconnection.

The issues were divided into four categories, and working groups established within the workshops to address each area. The first working group addressed the policies themselves. Working with the members of the FRICC, the initial statements

of agency policies were refined so that the rest of the workshop attendees could better understand the desired and required policies. The second working group addressed issues associated with access control to network resources. The third working group addressed the techniques required to support the sharing of networking resources in accordance with agreed upon policies. The fourth working group focussed on the end-to-end services required to support an interconnected set of networks.

Each of the working groups prepared summary reports of their deliberations. These reports are contained in Sections 3-6 of this document. The report of the policy working group attempts to summarize the existing policies of each of the agencies, particularly with respect to interconnection with other networks. The other three working groups focussed on the technology issues needed to be addressed in light of those policies. In each case, the working group report discusses the issues and develops an evolutionary capability with the goal of fully addressing the agency policies. Summaries of these reports are contained in the next section.

It is hoped that the results documented in this report will help the FRICC and the rest of the research community in achieving this exciting objective: a national research networking capability.

2. Workshop Summary

Driving the workshop were the policies of the individual agencies and a desire to interconnect the networks in a way that was satisfactory to those agencies. A prime policy driver appeared to be OMB Circular A130, which states that appropriate mechanisms must be used to assure some level of accounting for the use of the various networks. Another important policy driver was the need for agencies to assure that sharing of networks did not adversely impact the support of the individual agency users on their specific networks. This led in some cases to the need to be able to dedicate a portion (sometimes all during a specified time period) of an agency network to supporting its own users. Finally, the need to provide appropriate supporting end-to-end services, including security issues, led to the need for coordinating such services.

To facilitate the discussion of the technology issues and the presentation of results, it was decided to describe the evolution of capability in four phases. Phase 0 represented currently deployed and available capability. While not necessarily being currently used for the support of the policy issues, the capabilities of Phase 0 were viewed as being currently available and could be used starting today. Phase 1 consisted of capabilities that were developed and deployed at a limited number of sites. Thus, the issues involved in using such capabilities involved mainly those of widespread deployment (plus perhaps some limited amount of development associated with, e.g., porting of software). Phase 2 represented capabilities that were relatively well understood (little research required) but would require development activity before they could be used to support the policies for interconnection. Phase 3 capabilities require research to achieve, and thus represent the most future capability.

While these phases of capability represent evolution in availability, they should not be viewed as evolution in starting time for action. In all cases, research and development activities would have to start today in order that these capabilities be available in a timely manner.

As the working group on access control discussed the required technologies and mechanisms, it became clear that an important technology driver was the need to label packets with the appropriate information to make determinations of routing and resource allocation internal to the interconnected networks. For example, if certain links in a NASA network was to be restricted to use only by NASA users (even if accessing the network through an NSF network), it would be necessary to provide such labelling information in the packet. The report of the working group discusses the information that needs to be carried in such labels, requirements for authentication, and some potential experiments and development that should be carried out to achieve the required capability.

The working group on resource sharing focussed on the technologies that would allow fair sharing of resources between the participating agencies. The key issue that emerged from the discussions of this working group was the need to develop global

algorithms that permitted sharing and prioritization of the use of resources. As an example, it is relatively easy for an agency to block low-priority traffic from traversing its network during a period of high internal requirement. It is not so easy to do so and assure that the external users still can receive the resources they need from the interconnected internet.

The working group on end-to-end services focussed on those services that are required from a user's perspective from the overall system, and need to be coordinated across the interconnected networks. For example, directory and security services must be provided across the interconnected system. The key element emerging from the group discussions was the need to establish a consistent set of mechanisms to interconnect the various end-to-end services. These must be provided in a secure manner to assure that the security services fulfill their function.

The working groups identified the need to carry out supporting experiments and analysis to carry forward the interconnection of the networks, e.g., to make decisions about the need for stream versus transaction support. Each group developed a set of possible experiments and activities in accordance with the phases of development discussed above. These are summarized in Tables I-III.

A number of possible follow-on activities were identified to be passed on to the various Task Forces of the IAB. These are shown in Table IV.

In summary, the workshop identified a number of critical issues and identified areas where further research and experimentation is required. It is hoped that these results help provide a ''road map'' for how to satisfy agency policies and requirements in the interconnection of networks.

Table I
Access Control Projects

Phase 0	Access Control based on source/destination access matrix (for traffic not transiting network)
Phase 1	*Statspy* experiment to determine and define requirement for transactions
	‘‘ESnet hack’’ for limited access control based on source/destination addresses.
	‘‘Xerox hack’’ for limited access control based on source/destination addresses.
Phase 2	Coloring of stream packets
	Simple colors/labelling
	Route filtering for access control using source/destination addresses
	Incorporate ‘‘Xerox hack’’ into other gateways
	Authentication and signature architecture
Phase 3	Use of complex credentials
	Use of policy gateways in route computation

Table II
Resource Sharing Projects

Phase 0	Simple route filtering
Phase 1	Run *Statspy* to determine source/destination traffic flows (to comply with A130 traffic monitoring requirements)
Phase 2/3	50/50 resource management for link sharing
	Color packets and observe behavior to improve traffic monitoring
	Fast encryption of route and certificate packets, to secure traffic monitoring and control
	Fast mapping from source/destination to packet label/color
	Demonstration of gateway using soft state
	Define and support policy source routing
	Synthesis of source route
	Management controls and protocols
	Composition of policy terms
	Define and structure route set-up protocols

Table III
End-to-End Services Projects

Phase 0	User/process authentication using passwords (origin authentication)
	Mail relays for both function and system isolation
	Name domains system for host name to address mapping
Phase 1	User/process authentication using challenge/response or some other protocol (origin authentication)
	Secure-ID or other authentication technologies
	Challenge/response technologies (overlaps with the previous line)
	Kerberos (authentication server)
Phase 2	Authentication using certificates
	Integrity (MACs, checksums) and labelling
	Key distribution and management
	Secure mail (see RFC 1113)
	Certificates (see same RFC)
	Security of distributed white pages
	Integrity labelling, tools (MACs, checksums)
	Distributed white pages for the entire Internet
Phase 3	Use of VISAs
	Certification across peer domains
	Distributed computation
	National file system
	Trusted accounting
	Firewalls for end-to-end services
	Integrity of data across international boundaries with agreed upon cryptographic technologies
	Use zero-sum knowledge to have a third party to assure integrity without secrecy for such cases

Table IV
Projects for IAB Task Forces

ETETF	Handling of quality of service in gateways
ANTF	Phases 2 and 3 of resource sharing activities
IETF	Policy routing
Privacy	End-to-end privacy services
???	End-to-end services

3. Working Group on Interconnection Policies

Working Group 0 Members

Steve Wolff (Chair)	NSF
Guy Almes	Rice
Matt Bishop	Dartmouth
Brian Boesch	DARPA
Scott Brim	Cornell
Phill Gross	NRI
Dan Hitchcock	DoE
Russ Mundy	DCA
Tony Villasenor	NASA

Network resource sharing is encouraged by the potential for economies of scale both in communication link acquisition cost and in provision of value-added network services (the latter not yet demonstrated in the Internet, but consistent with telephone company experience); it is suggested by the Congressionally-ordered network study that resulted in the OSTP report *A Research and Development Strategy for High Performance Computing;* and it is mandated by OMB Circular A-130. Technical forces in the same direction include the additional connectivity each agency provides to its clients (actual or potential) by acquiring the use of nets belonging to other agencies at little or no additional cost, and the robustness afforded by the sharing of redundant paths or other forms of ‘‘excess’’ capacity.

The agencies represented on the FRICC, however, have differing missions and requirements, and these differences are reflected in differing rules and procedures for network usage. WG0 was created to explicate the rules for network use of the FRICC agencies, for those rules -- particularly the differences among them -- form the foundation upon which the technical specifications of ‘‘policy-based routing’’ must be built. This report, therefore, is the primary input to the technical Working Groups WG1, WG2, and WG3.

Making all FRICC agencies’ network use rules the same is NOT a goal of WG0. Each FRICC agency has more-or-less well-formulated rules for the use of its network in the absence of explicit interconnection with other networks and the attendant ‘‘foreign’’ traffic. These rules are given below. Currently, no agency has rules for interconnection with:

- networks of other FRICC agencies,
- networks of other countries,
- commercial networks, or
- ‘‘sensitive’’ networks (e.g., SDInet, NASA mission-critical nets);

consistent formulation of such rules will be discussed in future FRICC meetings.

It was however noted that, in dealing with subordinate (not peer) networks, NSF has required traffic presented to the NSFnet backbone to conform to NSF rules of acceptable use; DoE on the other hand is tending to the more liberal policy of carrying any traffic that meets the rules for acceptable use of the agency network offering the traffic.

3.1. Existing Policies, Summarized

The following is a summary of the existing policies for network usage of the FRICC member agencies.

NSF (draft, summarized):

- Purpose is to support scientific research and other scholarly activities.
- Use to support research or instruction at not-for-profit institutions of instruction and/or research is acceptable, whether all parties to the use are located or employed at such institutions or not.
- Activities in direct support of acceptable use are acceptable.
- Use for research or instruction by for-profit institutions may or may not be acceptable, and will be reviewed case-by-case.
- Commercial use by for-profit institutions is generally not acceptable.

DoE (draft, summarized):

- Use in which at least one party is supported by Energy Sciences funds is acceptable.
- Use by persons at DoE sites is acceptable, even if they are not supported by Energy Sciences funds.
- Advertising or promotional activities are not acceptable.
- Use in direct competition with commercial services is not acceptable.

NASA (draft, summarized):

- Purposes are to support NASA space science programs, to support collaborating science activities (e.g., with ESA, NOAA, USGS), and to support NASA contractors (e.g., those involved in building scientific sensors and spaceborne hardware).
- Other activities may be supported on a case-by-case basis, provided there is no impact to the NASA programs.
- No Eastern bloc access.
- Shared use of network facilities must be controllable and annually accounted for.
- NASA networking facilities may be made available for other uses and users on a cost-reimbursable basis.

- Direct competition with commercial services is not acceptable.

DARPA:

- Purpose is to support network research and other DARPA research objectives.
- There may be ‘‘forbidden routes’’ for some traffic.

DDN (excluding ARPANET and the proposed DRI):

- Use is for DoD business only, unless otherwise approved by JCS.
- All connections to other nets strictly regulated by mailbridges (now) or trusted guard gateways (future).
- Facilities must comply with DoD Security Architecture and with DoD Directive 5200.28 which requires C2 certification for sensitive unclassified information.

3.2. Refined Policy Statements

As a result of the first workshop discussions on policy, Dr. Cerf met with the various agency representatives to refine the policy statements. The results of these meetings were as follows. Note that these statements are those of the workshop and do not represent official agency policies. Each policy is represented in Clark’s Policy Term (PT) notation[1] and then described in English. The standard Clark Form for PTs (Hsrc,ARsrc,ARent)(Hdst,ARdst,ARexit){UCI}{Cg} FRICC={DOE,NASA,DCA,NSF} where H=Host, AR=Autonomous Region, src=source, dst=destination, ent=entry (previous hope), exit=exit (last hop, F=Federal Agency Net, Re=Regional, U=University, Co=Commercial Corporation, and Cc=Commercial Carrier. All PTs are a